

THE DEEP SPACE 1 AUTONOMOUS NAVIGATION SYSTEM: A POST-FLIGHT ANALYSIS

S. Bhaskaran, J. E. Riedel, S. P. Synnott, T. C. Wang

Navigation and Flight Mechanics Section, Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Abstract

NASA's New Millennium Program consists of a series of missions whose primary purpose is to demonstrate the feasibility of new technologies for spaceflight. Deep Space 1 is the first in this series of missions. It was launched on October 24, 1999 and has completed its first leg of the mission – flyby of the asteroid Braille – on July 29, 1999. An additional encounter is planned with the short period comet Borrelly in September 2001. The new technologies being demonstrated on DS1 include, among others, an ion propulsion system to provide maneuvering thrust, a combined visible/infrared/ultraviolet imaging instrument, and an autonomous navigation system. The purpose of this paper is to describe the computational elements of the autonomous navigation system and assess its performance in guiding the spacecraft to its first target. Some of the difficulties encountered during this leg, and how they were overcome, will also be described.

Introduction

For nearly 40 years, the primary method of navigating interplanetary spacecraft has been through the use of radiometric (two-way coherent Doppler and ranging) data, obtained by tracking the spacecraft using antennas at JPL's Deep Space Network (DSN) tracking stations, augmented by optical data from an onboard camera during encounters. This combination of data is very accurate, and has been used successfully to navigate to all planets in the solar system except Pluto, and several asteroids. This method of navigation does however require extensive use of a limited resource (the DSN), as

well as ground personnel needed for analysis. For many years, it has been known that, in principle, the navigation system could be fully automated and self-contained by using a camera taking triangulation images of solar system bodies to determine the spacecraft's position and velocity, and then computing and executing maneuvers onboard to deliver the spacecraft to its target. However, the inherent risk in trying this for the first time on a science mission has precluded an actual flight test. NASA's New Millennium Program fills this void by funding a series of missions whose primary purpose is to demonstrate the feasibility of new technologies for spaceflight. The asteroid/comet flyby mission named Deep Space 1 (DS1) was the first in this series of missions. In addition to autonomous navigation, other new technologies tested included an ion propulsion system to provide maneuvering thrust and a combined visible/infrared/ultraviolet low mass and power imaging system named MICAS (Miniature Integrated Camera and Spectrometer).

Although the principles behind an autonomous navigation system are fairly simple, the implementation posed several challenges when applied to the flight of an actual spacecraft. The purpose of this paper is to describe the computational algorithms used by the autonomous navigation system (autonav) and how they performed in guiding the spacecraft to its first target.

The Mission

DS1 was launched on a Delta 7326 rocket on October 24, 1998, and flew by its first target, the asteroid Braille, on July 29, 1999. Although originally planned to flyby two more targets, the loss of the onboard star tracker for attitude control in November, 1999 prevented the spacecraft from performing its nominal thrust profile using the ion engines to provide the energy needed to reach both of these

Copyright ©2000 by the American Institute of Aeronautics and Astronautics, Inc. The U. S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for government purposes. All other rights are reserved by the copyright owner.

targets. As of this writing, the spacecraft team is working on software to use the MICAS imaging camera to replace the star tracker, with the plan being to start another thrust profile in the summer of 2000 to flyby one of the targets, the comet Borrelly, in September 2001.

From an autonav standpoint, the two other technologies having the largest impact on the design of the autonav system are the use of ion propulsion engines, and the characteristics of the MICAS camera. Unlike chemical propulsion systems which burn for short periods of time at very high thrust, the IPS produces very little thrust but is capable of burning for very long periods of time. Ionized xenon is accelerated by passing it through a charged grid before exiting out of the nozzle. The resulting thrust is on the order of millinewtons, with specific impulses reaching values in the thousands of seconds (as compared to 200-400 seconds for chemical rockets). The thrust can be throttled by varying the voltage on the grids; for DS1, the IPS has about 100 throttle levels, with a thrust range of 20 to 90 mN. Since the power is generated from the solar arrays, the maximum achievable thrust depends on the distance to the sun.

The characteristics of an IPS trajectory are different from those using chemical engines. Trajectories using chemical engines have long coast periods punctuated by near-instantaneous velocity changes at given times to achieve course corrections. IPS trajectories, on the other hand, are characterized by long thrusting periods of weeks to months, interspersed with coast arcs when the IPS is shut off. For DS1, the thrusting periods have the dual purpose of providing enough energy to the spacecraft to reach its targets, and correcting launch injection, orbit determination, and maneuver execution errors to achieve the desired targeting conditions. One of the major design considerations for the autonav system was the process of computing alterations to the nominal IPS thrust profile and corrective IPS burns for targeting.

The MICAS camera has four channels, two in the visible light spectrum, and an infrared and an ultraviolet spectrometer; only the visible light channels were used by autonav. Of these, the primary one used during the majority of cruise and approach to the asteroid was a standard Charge-Coupled-Device (CCD) chip with a 1024 square pixel array. Each pixel has a field-of-view (FOV) of about $13 \mu\text{rad}$ for a total FOV in the CCD of 1.3 mrad , or 0.76° . The other visible channel is an experimental Active Pixel Sensor (APS) array which has a 256 square pixel array and a total FOV of 4.6 mrad (0.26°).

Both were coupled to a telescope with a focal length of 685 mm whose boresight is fixed to the spacecraft (thus, the entire spacecraft has to be slewed to point at particular region of the sky). Also, both have 12 bit digitization, resulting in data numbers (DN) values for each pixel ranging between 0 (no signal) and 4095 (saturation). Prior to launch, the MICAS development team had determined that the CCD would have excessive charge bleeding when exposed to a bright, extended object such as during the flyby period. Thus, it was decided to use the far less sensitive APS in the final 20 minutes or so of terminal tracking during flyby. This decision had a major impact on the success of the flyby tracking, as will be described later.

Elements of the Autonav System

The entire autonav code was designed to be as self-contained and modular as possible to make it adaptable to other missions as well as for DS1. The system can be divided into three components; an executive which interacts with the rest of the flight software and is responsible for scheduling and executing events, a real-time ephemeris server which provides spacecraft and solar system body ephemeris information to other flight software elements on a quick turnaround basis, and the computational elements which perform the fundamental navigation updates. The first two components will not be described in detail here; more information can be found elsewhere¹. This paper is primarily concerned with the performance of the computational elements, which include the following functions: 1) orbit determination, 2) maneuver planning, and 3) encounter target tracking. Each of these will now be described in detail.

Orbit Determination

Orbit determination is the process by which the spacecraft's state (position and velocity) and other parameters relevant to the trajectory, such as non-gravitational accelerations acting on the spacecraft, are estimated. In order to keep this process as self-contained onboard the spacecraft as possible, the only data used to obtain an OD solution are images taken of solar system bodies (asteroids in this case) by the MICAS camera. The principle is the following; each sighting of an asteroid in the camera FOV places the spacecraft in a position along that line-of-sight (LOS). Two or more such sightings of different asteroids fixes the spacecraft's three-dimensional position by triangulation. The stars in the background are needed to determine the inertial pointing di-

rection of the camera boresight (since the stars are so distant, their inertial directions will not change measurably when seen from different locations in the solar system, so they can be thought of as "fixed" in the sky). In practice, however, two simultaneous sightings are not practical with one camera, and instead, a series of LOS fixes are taken of several asteroids. Several clusters of sightings are then incorporated into a least-squares filter to obtain an OD solution. The accuracy of this type of data is dependent on several factors, including the angular separation, brightness, and distance to the imaged asteroids, the resolution of the camera, the ability to pinpoint the location of the asteroid in the camera frame (centerfinding), the accuracy of the camera pointing information, and the knowledge of the asteroid ephemerides. Pre-flight analysis and simulations had shown that the designed DS1 system was capable of meeting the navigation requirements for asteroid/comet flybys². For clarity in the following descriptions, the term "beacons" is used to denote the asteroids used solely for triangulation, while "target" refers to the objects being encountered (asteroid Braille for the first leg of the mission).

Image processing is the first step in the OD process. Its primary purpose is to predict the locations of beacons and surrounding stars at given times, determine the center of the asteroid and the stars in the camera frame, and compute the associated pointing of the camera boresight. The ability of the navigation system to perform autonomously hinges on its ability to accurately perform the centerfinding and ensuring that bad data do not corrupt the solution.

Computing predicts of beacon asteroids is the simplest of these procedures. A list of beacon asteroids to observe as a function of time for the relevant portions of the mission is stored onboard the spacecraft, along with ephemerides of all the beacons. At predetermined times, the current spacecraft trajectory is differenced with the nominal ephemeris of given beacon to get the relative pointing vector. This information is then passed to the spacecraft attitude control system (ACS) which slews the spacecraft to the correct orientation at the correct time and shutters the picture with the provided exposure length. After the frame is taken, it is stored as a file in the main flight computer which can be accessed by autonav.

The image processing itself has two stages; the first performs a coarse registration of the beacon and stars in the frame, and the second fine tunes the center locations to subpixel accuracy. Originally, the coarse registration process involved locating all the bright spots in the frame above a give thresh-

old and then matching the pattern of spots to the known pattern of stars that was expected. Due to problems encountered with MICAS however, a second algorithm was loaded onboard mid-flight which relies on the relatively good a priori knowledge of the inertial pointing direction provided by the ACS star tracker. Around a 50-100 pixel predicted location of each star, the center-of-brightness (COB) is computed. All the predicted locations for the stars are then be shifted by the difference between the COB and the nominal location of the particular star, and the integrated DN value in the regions surrounding the updated positions for all the stars is computed. This process is repeated for each star in the frame; the shift providing the maximum integrated DN value, summed over all the stars, is the offset applied to the predicted locations for the coarse registration.

The second step to fine tune the center locations uses a cross-correlation technique inherited from a similar method used on the Galileo mission. Because the spacecraft attitude wanders slightly between its commanded deadband during the exposure time, the image of the stars and beacon are smeared in a complex, and unpredictable, pattern in the frame. Thus, a simple Gaussian shape cannot be used to locate the centroid of each object; instead, the complicated shape itself is exploited. The procedure is to use the image of every object in a frame as a template to cross-correlate with all the other objects. Thus, every object in the frame will have associated with it a set of pixel/line shifts which obtain the peak correlation of its shape with that of the others. This information can be combined in a least-squares process to find an ensemble set of shifts for all the objects, holding one of the objects (usually the beacon) fixed, which best describes the locations where the signals have peak responses with respect to each other³. A post-processor then screens the results for low signal and bad data by deleting objects which do not pass a threshold response. From prior experience with the Galileo mission and from tests done on images taken from a ground telescope, pre-flight estimates on the accuracy of this procedure was around 0.1 pixels. For reasons described later, the actual performance in flight varied between 0.2 and 0.8 pixels.

Once the centroids are determined, the inertial pointing direction of the camera boresight can be computed from the star locations using the ACS supplied values as a starting guess. If two or more stars are available, then all three components of the pointing, right ascension (RA), declination (DEC), and the twist (TW) around the boresight, can be determined to better than 10 rad. If only one star

is available, then the TW knowledge is degraded. If no stars are available, then the image was discarded during cruise. The updated pointing information, along with the observed centroid of the beacon, is written to a file for use by the filtering link in the OD process.

The dynamical model of the spacecraft trajectory employed by autonav include, in addition to the central body acceleration, third body perturbations from other planets, solar radiation pressure, thrust forces from the IPS, and a general bias acceleration term used for additional but small unmodelled forces. The propagation of the spacecraft motion was done entirely in a heliocentric, Earth Mean Equator of 2000 coordinate system using a 7-8th order Runge-Kutta numerical integrator. Third body perturbations included those from Venus, Earth, Moon, Mars, Jupiter, and Saturn; errors caused by the exclusion of the remaining planets were well within the noise of the data. For solar radiation pressure, the spacecraft presents a large cross sectional area from the solar panels which maintain, for the vast majority of the cruise, a fixed orientation relative to the sun. Thus, for simplicity, the model of solar radiation pressure assumed a spherical shape for the spacecraft whose area was equal to the area of the panels. The unmodelled acceleration term was a bias in three cartesian axes applied over the whole data arc. Nominally, this term is set to zero but is estimated in the filter to absorb residual accelerations not accounted for in the rest of the dynamic model.

The treatment of thrusting events, both from the Reaction Control System (RCS) and IPS thrusters, deserves some discussion. RCS thrusts are produced by several hydrazine engines placed around the spacecraft and are used primarily to control the spacecraft attitude, and to provide relatively small thrusts for course corrections. The spacecraft records thrusting events from both the IPS and RCS by passing a message to the autonav system every time the thrusters are turned on. This information comes at a fairly high rate (about 1 Hz); autonav compresses the information by averaging over varying time intervals and writes it to a special history of events file. This history file models RCS events as discrete velocity changes (ΔV) at specified times, and IPS events as a piecewise constant acceleration with a first order polynomial for the direction. For integrating through past events from some epoch to the current time then, the integrator reads the history file and adds in the effects of the ΔV s and the IPS accelerations along with the other dynamic forces acting on the spacecraft. Be-

cause the thrust event values provided to autonav are based on mathematical models of RCS and IPS engine performances, their accuracies varied. During flight, it was found that the RCS ΔV magnitudes computed onboard were underestimated by a factor of nearly 2; the model parameters were only updated many months after launch. The IPS information, on the other hand, was calibrated to good precision and was accurate to the 1-2% level. Furthermore, the filter estimates a scale factor which multiplies the thrust magnitude to account for deficiencies in the calibration.

For propagating the trajectory into the future, the thrust is handled differently. The autonav system carries onboard a maneuver events file, separate from the history file, which has on it all the deterministic IPS and RCS events predicted for the future. The information on this file is originally supplied from ground trajectory planning, but the autonav system overwrites it every time it computes a retargeting maneuver. The intent of this file was to provide information about IPS and RCS targeting and course correction maneuvers only, but it was found in flight that attitude control RCS events caused drifts of the trajectory over the long term which, if not accounted for, caused unacceptable errors in targeting. This was primarily a problem in the months prior to encounter, so the maneuver file during this time carried predictions of the effects of small attitude control thruster firings as discrete events at specified times.

In general, the problem of computing an orbit from the optical tracking data is a nonlinear one; however, the autonav system uses the standard practice of linearizing about a nominal trajectory and estimating corrections to the nominal state parameters which minimize the data residuals in a least-squares sense⁴. The estimated parameters include corrections to the spacecraft state (position and velocity), the bias accelerations, and the IPS thrust scale factors. The technique used is the batch epoch-state filter whereby all the data in a given batch are used to estimate the state at the epoch time of the batch, after which the solution and associated covariance are propagated through the data arc to the current time. The normal equations produced by the least-squares batch method are solved for using the U-D factorization method⁵. The batch length for computing solutions varied, but in general, an attempt was made to keep the arc to close to 30 days. This was found from preflight studies to be an optimal length; shorter arcs do not provide enough data to get good state estimates while for longer arcs, the addition of new data does not improve the

epoch state solution due to the corruptive nature of the non-gravitational accelerations (primarily from thrust events)². The a priori covariance on the estimate was kept large to prevent filter divergence, so information from a previous batch solution was not explicitly carried over. However, an OD update was performed fairly frequently with constant arc lengths. When doing weekly updates for example, the earliest weeks data would be dropped and the latest weeks data added. Thus, the solutions overlap, mitigating the effects of losing information with an unconstrained epoch state.

A brief description of the OD process is as follows. At roughly weekly intervals, the spacecraft sets aside a 4-6 hour block of time for beacon asteroid observations. Using a list of pre-planned beacons and its onboard stored spacecraft and asteroid ephemerides, the spacecraft turns and shoots a series of frames of each beacon. After each frame was shuttered, the frame is immediately processed by the image processing link, and the pertinent information appended to a file of previously stored processed data. After all the frames were taken and processed, the OD link computes updates to the trajectory using all the data in the arc; the solution is then be mapped to the current time and into the future. The new trajectory overwrites the old trajectory information on the spacecraft ephemeris file, and becomes the nominal one until the next OD solution. Thus, the spacecraft has a constantly updated knowledge of its own whereabouts which are used by other elements of the flight software, notably ACS. Various cleanup tasks are also performed, such as truncating the history and/or image data files. Finally, if called for, a maneuver targeting computation is performed.

Maneuver Planning

The purpose of the maneuver planner is to compute the course corrections needed to achieve the target flyby conditions. The course correction can be implemented in one of three ways: changing the magnitude, direction, and duration of the nominal thrust profile that the ion engine flies, adding a Trajectory Correction Maneuver (TCM) using the ion engine, or adding a TCM using the hydrazine attitude control thrusters. In all three cases, the correction maneuver is computed by taking partial derivatives of the desired target condition with respect to the control parameters, inverting this matrix, and multiplying by the residual error formed by the difference between the current and desired target condition. Thus, it is a linear control process and requires a reasonably good starting condition in order to con-

verge. The initial design of the nominal trajectory is done on the ground and uplinked to the spacecraft. More details of the maneuver design process can be found elsewhere⁶.

Encounter Target Tracking

The complete set of dynamics and filter described above was used as the primary mode of operation for OD solutions throughout cruise and until several hours prior to encounter. At this stage, no more maneuvers are performed and the main purpose of autonav is to maintain visual lock on the target as it flies by. This requires rapid updates of the state as tracking images are taken, and the OD link is not fast enough for this purpose. For this reason, a compact filter, termed Reduced State Encounter Navigation (RSEN) was adapted from a similar one used on the STARDUST mission⁷. RSEN uses the final position and velocity from the main OD link at 30 minutes prior to encounter and models the trajectory as a target-centered straight line through encounter. As frames are taken and processed, the epoch position alone is updated and a new linear course computed. The spacecraft relative target position and updated time of closest approach are passed to the ACS system to keep the target in the camera FOV.

The image processing for RSEN differs from the mainline OD in three important aspects. First, the smaller FOV and less sensitive APS camera alone is used; no CCD images are processed by RSEN. Secondly, since the target is an extended object at this close range, the image processing consists solely of computing a brightness centroid. The region to search for the brightness centroid in the frame is limited by mapping the covariance of the state into the camera FOV and using only the pixels which fall within a 3σ box surrounding the predicted location. Finally, since no stars are visible, the pointing information is provided by the ACS star tracker with no improvements. This corrupts the estimated state to some degree, but analysis had indicated that the accuracy is still sufficient to track.

Operational Results

MICAS Problems

Prior to launch, all the computational elements (except RSEN) of the autonav system were extensively tested by Monte Carlo simulations using realistic models for the performance of the engine and the spacecraft ACS system that were known at the time. These tests indicated that autonav was capable of

delivering the spacecraft to its close flyby of an asteroid without need of ground intervention and within reasonable margins for propellant usage⁵. The actual flight, however, turned up several unanticipated problems, largely due to the performance of the MICAS camera, which made autonav extremely difficult to accomplish. The earliest indication of trouble was when the first images from MICAS were taken and downlinked on November 6, 1998 (Figure 1). These images showed that the frame was severely corrupted by stray light, clearly visible in the left portion of the frame, which has a background brightness level of over 1000 DN from the right portion. It was quickly apparent from these images that the original software loaded onto the spacecraft to perform the initial star registration would not work due the large variation in the background DN levels caused by the stray light. Additional images taken on November 17 further showed that the intensity of the stray light seemed to vary as a function of the angle between the camera boresight direction and the direction to the sun (hereafter referred to as the cone angle). Intensive efforts were then undertaken to characterize the problem during the early months of 1999 by taking images at various cone angles. It was found that there were two regimes of stray light problems; between cone angles of 110° to 180° , the light was in the form of a “card” as seen in the first frames which increased in intensity as a function of cone angle, and at cone angles less than 110° , a “blowtorch” of light in the upper third of the frame appeared (Figure 1). The former problem is more benign because the image is still usable at long exposure durations; the latter more serious because the blowtorch saturates very quickly and bleeds onto the remainder of the frame at exposures longer than about 10 seconds.

Another unanticipated problem surfaced in early December of 1998 when a set of 27 images of a dense star cluster were taken and downlinked. The purpose of the images was to characterize the geometric distortions in the camera focal plane which, if not modelled properly, would cause errors in the orbit solutions. It was found from processing these images that the standard model of the distortion field that the flight software uses did not accurately represent the high frequency variations. In particular, it was found that fitted parameters to the model computed from the locations of the star image centroids resulted in residual rms errors of nearly 1 pixel. This would cause, for example, a triangulation sighting of an asteroid which was 1 AU distant, a spacecraft position fix error of about 2000 km – much larger than desirable. For comparison, the same model used on

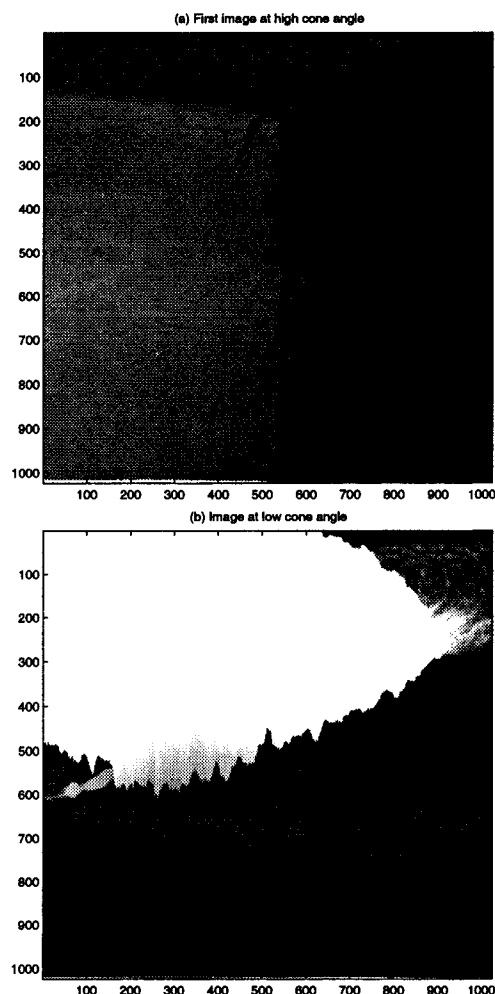


Figure 1: Examples of images at high and low cone angles

the Voyager and Galileo missions had rms errors of less than 0.1 pixel. Thus, in addition to handle the stray light, another software fix to account for the unusual distortions had to be developed as well.

A third major problem with MICAS that affected autonav was the reduced sensitivity of the camera. Originally, the preflight specifications called for the ability to image stars and asteroids with visual magnitudes of 11.0 or better at reasonable exposure durations. This allowed the use of up to 80 asteroids for triangulation, and provided a good geometric spread of beacons for any given OD opportunity which reduces the effects of asteroid ephemeris and other systematic, as well as random, errors. This requirement also ensured that several stars are in the FOV for determining the absolute pointing direction. In practice, however, the camera sensitivity turned out to be less than anticipated, with the result that only objects whose magnitudes were brighter than 9.0 9.5 were reliably detected. This severely constrained our choices for beacons; only the larger and brighter asteroids were now available and the system had to cope more often with frames which had less than two stars. As a consequence, the original charter of autonomy was descoped, with the autonav team on the ground taking responsibility for carefully selecting the beacons which had the proper brightness and sufficient stars in the background.

For these and other reasons, the initial use of the computational elements of the autonav system, originally planned for early December, was delayed while new software was developed and tested (other autonav components which were not affected by the camera problems were started shortly after launch and performed nominally). In the meantime, images were still being taken by the spacecraft and downlinked to the ground for analysis.

OD Results During Interplanetary Cruise

During flight, two opportunities for uploading upgrades to the autonav software were available; the first was on February 8, 1999, and the second on June 16, 1999 (dubbed the "M4" and "M6" loads, respectively). The M4 load included new algorithms for initial registration and updated parameters for the original distortions model. M6 included the new distortion model and its associated parameters, the RSEN tracking code, and new code to perform frame differencing, whereby a "background" frame is first taken and differenced with the normal beacon frame. This capability was added because it was found that the stray light was fairly stable over small differences in cone angle, and the differencing removed

most of the large wavelength variations from the stray light. For evaluating the performance of the onboard autonav, many of the raw images, as well as centroiding results, were downlinked for analysis. Also, since standard radio navigation was concurrently being performed on the ground, the results from the onboard OD could be compared with the radio results (which is accurate to better than 100 km in heliocentric space).

On March 8, the spacecraft computed its first completely autonomous onboard solution on using data starting on February 18. Subsequently, OD updates were performed at roughly weekly intervals with little ground intervention until early June, 1999, with an enforced gap in mid-May due to testing of other experimental software components. Figure 2 plots the onboard OD accuracies during the M4 time frame as the difference of the autonav estimated positions and velocities with the ground radio results. For simplicity, only the rss of the three position and velocity components are shown. Also plotted are the 1σ uncertainties in the position and velocity. The plot shows that after an initial position error of about 4000 km, the discrepancy ballooned to over 14,000 km before settling back down to the 4000 km range, and finally ending up less than 2000 km. The cause for the large jump was found to be an erroneous parameter describing the alignment of the camera boresight that was sent to the spacecraft; once this was fixed, it took several weeks before the data using the incorrect parameter exited the data arc. The marked improvement at the end was due to the addition of the asteroid Vesta, a large and bright asteroid, to the solution which provided better geometric spread than had been seen before (this can be seen in the position sigma, which dropped to about 600 km where before it had been hovering in the 1600 km range). The velocity errors varied between about 1.5 and 7.0 m/s. Both position and velocity errors were about an order of magnitude larger than preflight estimates, and usually several factors larger than the formal uncertainties. This can be primarily ascribed to the loss of beacons and the lower fidelity distortion model, but another factor was the fact that the onboard estimates of RCS thruster firings provided by ACS were found to be off by nearly a factor of 2. The latter was corrected by M6.

Figure 3 plots the onboard OD errors and associated sigmas computed after the M6 software load. Following the initial position error of around 1200 km, the results increased to 3400 km in mid-June before settling down to a 700-1000 km range. The sudden increase was, once again, due to erroneous parameters that were uploaded, and a corrected file

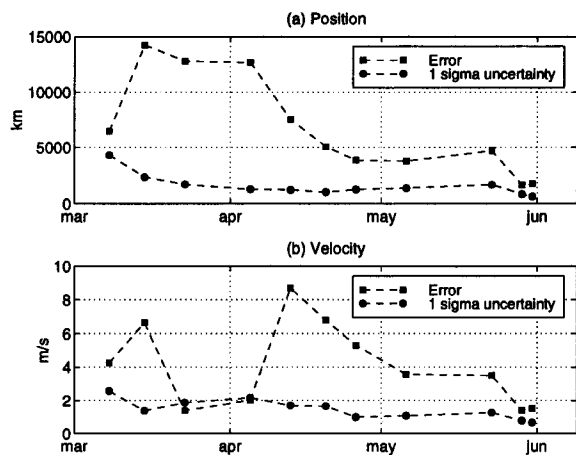


Figure 2: Position and velocity errors and uncertainties during the M4 time period.

was sent up in mid-June. The steady state velocity errors during this time period was less than 0.5 m/s, and dropped as low as 0.2 m/s in parts of the arc. Clearly, the improvements in software and beacon selections had an effect on the OD accuracies, which were now only a factor of 2 or 3 larger than preflight estimates. Also note that with the M6 software, the errors are generally consistent with the formal uncertainties of the solutions. Overall, the system was now performing about as well as it could in theory given the geometrical constraints provided by the data. The loss of many beacons is reflected in the plots of the position uncertainties, which can vary markedly from one week to the next as opportune beacons rise and set. Had the camera response been as expected, the curve would have been much smoother and the solutions would not have been as sensitive to the gain or loss of any single beacon.

Between the upload of the corrected parameter file in mid-June and the latter part of July, orbit determination onboard the spacecraft was fairly independent of ground intervention. Due to a combination of the reduced geometric information and unmodelled high frequency camera distortions, however, the accuracy being obtained was not sufficient to support fully autonomous maneuver planning. In particular, the maneuver planner was set to execute TCMs at Encounter (E) - 20, 10, 5, and 2 days, but only if the discrepancy between the predicted and desired target conditions were larger than 2 sigma of the formal error. Further complicating the matter was the fact that the spacecraft performed a rehearsal of the encounter sequence in early July, which involved executing an actual maneuver based on simulated data for encounter. For these reasons, the TCMs

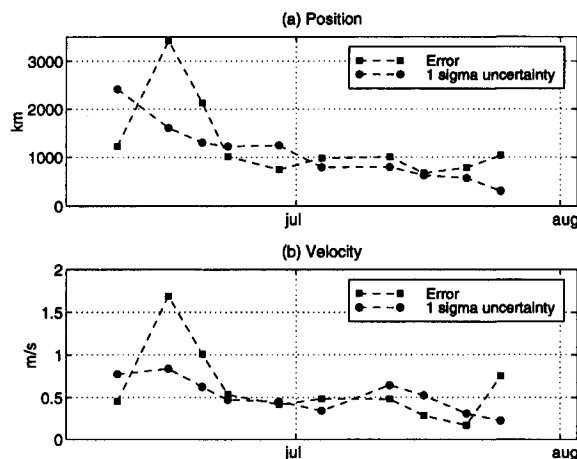


Figure 3: Position and velocity errors and uncertainties during the M6 time period.

computed onboard were checked on the ground to see its effect on the spacecraft trajectory based on the more accurate radio data. Using this information, TCMs at 20 and 2 days were cancelled. The E - 10 day TCM was executed by carefully planning the rehearsal TCM to be in a direction which would be advantageous. The E - 5 day TCM computation was based on ground radio data due to an autonav flight software error which occurred on July 21 and caused corruption of the autonav file system. The target condition in all cases is referred to in the "B-plane" coordinate system. This coordinate system is a plane centered on the target body and perpendicular to the incoming asymptote. The flyby aimpoint that the spacecraft was being targeted to is located in this plane at a horizontal location ($\mathbf{B} \cdot \mathbf{T}$) of 12 km, and a vertical location ($\mathbf{B} \cdot \mathbf{R}$) of 9 km.

Approach Results

Based on ground data for the size, shape, and albedo of Braille, it was predicted that Braille would be visible in the MICAS CCD about 3 - 5 days before encounter. In practice, the first sighting of Braille occurred at E - 3 days, but only after extensive image processing on the ground (the onboard software was unable to detect the signal of Braille at this time). By registering the absolute orientation of the frames using a nearby guide star and then co-adding all the downlinked frames taken on July 26, a faint signal appeared which was likely due to Braille. Subsequent images taken and downlinked on July 27 verified that the signal, though still quite dim, was indeed from Braille (Figure 4). A ground computed OD solution based on this data indicated

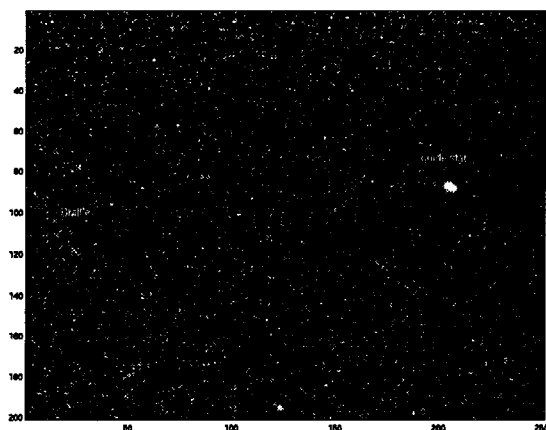


Figure 4: 1st unambiguous sighting of Braille in CCD.

an ephemeris error for Braille of about 350 km (representing about 2 sigma change from the nominal ephemeris), placing the spacecraft in the B-plane at a $B \cdot R$ of 248 km and $B \cdot T$ of 371 km, with an uncertainty of about 50 km (Figure 5). This was large enough to warrant a ground computation of a maneuver at E - 1.5 day to retarget the spacecraft, which was executed nominally.

Following the E - 1.5 day TCM, a set of 18 frames of Braille was taken on July 28 at 00:03:00 UTC (E - 28 hrs). The onboard software was still unable to locate Braille in any of these images due to its signal being below the detectability threshold. From analyzing 5 of these images on the ground though, the spacecraft's current flyby location was computed to be at a $B \cdot R$ of 7.6 km and $B \cdot T$ of 21 km, with an uncertainty of roughly 16 km (Figure 5). Due to the size of the uncertainty, a planned E - 18 hr maneuver was cancelled.

At 11:30 UTC on July 28, 18 more frames of Braille were taken. Inspection of spacecraft telemetry data downlinked much later indicated that autonav had finally locked onto signal from Braille. Unfortunately, during the OD processing, a latent bug in the autonav software caused a spacecraft safing event, wiping out the latest onboard trajectory result. Following recovery of the spacecraft from safing, three of the images were downlinked, and from this data, a ground OD solution computed. This solution showed that the spacecraft was now at a $B \cdot R$ of 4.0 km and a $B \cdot T$ of 6.0 km, with an uncertainty of around 11 km (Figure 6). This result was somewhat disconcerting in that it represented a greater than 1 sigma deviation from the previous

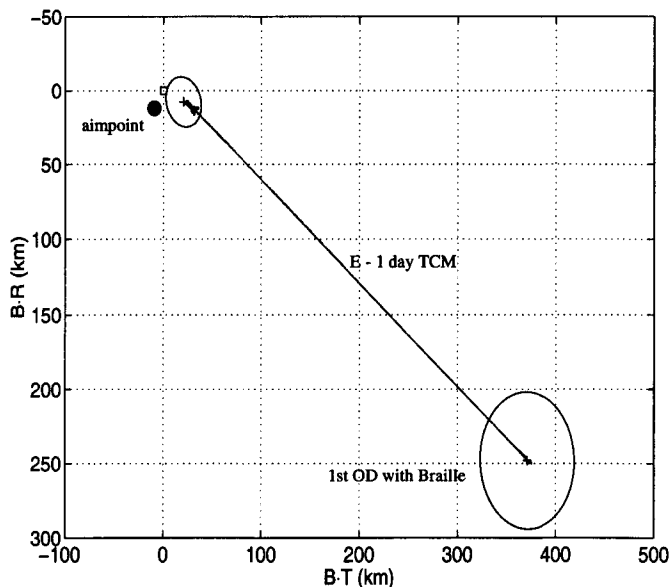


Figure 5: Post-Braille detection B-plane solutions.

solution and furthermore, that the spacecraft was uncomfortably close to an impact trajectory. Thus, a TCM was developed to retarget back to its nominal aimpoint and uplinked to execute at E - 6 hours. It should be noted that the turnaround time from receiving the images to developing a maneuver sequence to be uplinked took no more than 1 hour.

Following the TCM, four sets of Braille images were taken at E - 4, 3, 2, and 1.5 hours. These processed onboard normally, but did little to change the onboard estimate of the flyby trajectory, which, following the nominal execution of the 6 hr TCM, was computed to be at its targeted position. This information was handed over to the RSEN subsystem at E - 27 minutes.

RSEN Results

Real time Doppler data taken during encounter indicated that the flyby had proceeded safely and the spacecraft was operating normally. Following encounter, however, the downlinked close approach images did not have the asteroid in it, either in the CCD or APS frames, indicating a failure of RSEN to properly track the asteroid. After all the onboard files were played back, only two frames showed Braille; both were in the CCD and taken approximately 15 minutes after closest approach. This, along with telemetry data produced by RSEN and the ACS system, enabled the autonav team to reconstruct the events during encounter and pinpoint the probable cause of the failure.

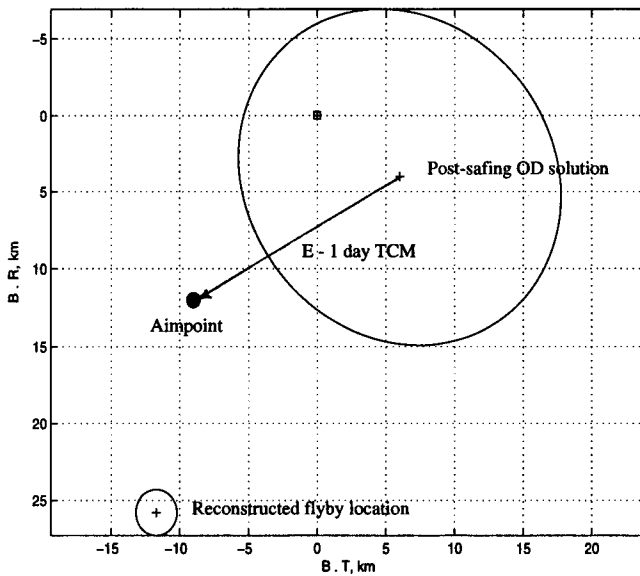


Figure 6: Post-safing B-plane solutions.

Using all pre-encounter images of Braille and the two post-encounter ones, the actual flyby location was computed to be at a $B \cdot R$ of 25.8 km and $B \cdot T$ of -11.7 km, with an uncertainty of 1.5 km. (Figure 6). This was slightly greater than 1 sigma of its targeted location, but still well within the ability of RSEN to track. The telemetry from RSEN showed that APS frames between E - 27 minutes and E - 24 minutes were being processed normally, but that the signal from Braille had not appeared above a preset threshold. At E - 24 minutes, a spurious signal did appear above the threshold and spoofed RSEN into making about a 30 km error in the B-plane estimate. This error biased subsequent commanded pointing directions to a level where Braille was no longer in the APS FOV. However, using ACS telemetry of pointing angles and the reconstructed trajectory, it was determined that Braille was still within the CCD FOV until E - 5 minutes, and then briefly again at E - 3 minutes. Unfortunately, due to onboard storage constraints, neither the science nor autonav teams planned CCD images to be taken and stored prior to E - 2 minutes. By E + 15 minutes, the spacecraft had reverted to its pre-RSEN trajectory, and captured Braille in the CCD. In addition, four APS images taken during this time also had Braille in its FOV, but only in one of them was there a signal of any note. This signal was barely above the background, and could only be identified because it matched the predicted location of Braille.

The question then remained as to why the signal from Braille was so weak. Data informally presented by various sources prior to encounter had led the autonav team to believe that Braille would produce a strong response in the APS with the selected exposure durations (about 5 seconds), and the threshold values for detection were set accordingly. The lower than expected signal in the post-encounter CCD frames, and the lack of signal in the APS, however, indicated that Braille was perhaps 40 times dimmer than predicted. Also, the unusual shape of Braille probably presented a projected area to the approaching spacecraft which decreased its integrated brightness. Finally, the response of the APS detector was nonlinear and poorly characterized, resulting in an optimistic expectation of its sensitivity.

A combination of the above factors led to autonav being unable to detect Braille in the E - 27 to E - 24 minute period. The noise spike which perturbed RSEN's estimate of the trajectory pointed to a design flaw in the software, which did not check for persistence of a signal before accepting it as real. This feature was not thought necessary on the assumption that the signal from the target would be so strong that the effects of noise would be negligible. This was verified in an experiment where the two CCD images of Braille, which had the level of signal that was expected in the APS, were used as data for RSEN. RSEN was able to pinpoint the flyby to within two km using just the two images, and the noise in these frames had no effect. However, even had RSEN had not been spoofed by the noise, the end result would have been the same because the knowledge of the trajectory prior to initiating RSEN was not sufficient to track the asteroid open loop beyond the E - 2.5 minute time period.

Summary and Conclusions

After the initial difficulties with parameters following the M4 and M6 uploads, autonav, at least as far as image processing and OD are concerned, was performing its task more or less without ground intervention. The maneuver computation routines were slightly less autonomous in that the state information it was given occasionally had to be uplinked from the ground, but it still worked as designed onboard. For the interplanetary cruise portion of the mission, the autonav system was deemed validated by the project, and indeed, following the Braille encounter, had been in control of the spacecraft with little intervention and no radio backup until the star tracker failure in November. The accuracies obtained by the M6 version of the software is more

than sufficient for this mission phase, and a planned M7 load would have improved on its overall reliability. During the approach phase to a target, however, the ground would have to intervene because the OD accuracy is not good enough to provide a robust sequence of maneuvers which would take the spacecraft progressively closer to its aimpoint. This failure, however, can be attributed almost entirely to the fact that the MICAS camera did not meet its design requirements. For future missions, it is clear from the DS1 experiment that, with a well designed and characterized camera, an autonav system could easily perform both mission phases reliably.

Following the lessons learned from the Braille encounter, the RSEN tracking code has been modified to make it more robust for the Borrelly encounter. These modifications, however, cannot be used as a substitute for proper characterization of the target being flown by. This is an important point to note for other missions which rely on an autonomous system for closed loop tracking of a small, relatively unknown object, such as the STARDUST mission which uses a very similar algorithm during its comet encounter. In this matter, DS1 has performed a valuable role in real-world testing of unproven techniques which can enhance the scientific returns from more conventional missions.

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

1. Riedel, J. E., S. Bhaskaran, S. P. Synnott, S. D. Desai, W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy, G. W. Null, W. M. Owen Jr., R. A. Werner, and B. G. Williams, "Navigation for the New Millennium: Autonomous Navigation for Deep-Space-1", Proceedings 12th International Symposium on Flight Dynamics, Darmstadt, Germany, June 2-6, 1997.
2. Bhaskaran, S., J. E. Riedel, S. P. Synnott, "Autonomous Optical Navigation for Interplanetary Missions", *Space Sciencecraft Control and Tracking in the New Millennium*, E. Kane Casani, Mark A. Vander Does, Editors, Proc. SPIE 2810, pp32-43, 1996.
3. Liebelt, P. B., *An Introduction to Optimal Estimation*, Addison-Wesley, 1967.
4. Bierman, G. J., "Measurement Updating Using the U-D Factorization", Proc. 1975 Conf. on Decision and Control, Houston, TX, pp. 337-446.
5. Bhaskaran, S., Desai, S. D., Dumont, P. J., Kennedy, B. M., Null, G. W., Owen, W. M., Riedel, J. E., Synnott, S. P., and Werner, R. A., "Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation System", AAS Paper 98-193, AAS/AIAA Spaceflight Mechanics Meeting, Monterrey, CA, February 1998.
6. Desai, S. D., Bhaskaran, S., Bollman, W. E., Halsell, C. A., Riedel, J. E., Synnott, S. P., "The DS-1 Autonomous Navigation System: Autonomous Control of Low-Thrust Propulsions Systems", AIAA Paper 97-38819, AIAA Guidance, Navigation and Control Conference, New Orleans, LA, August 1997.
7. Bhaskaran, S., Riedel, J. E., Synnott, S. P., "Autonomous Nucleus Tracking for Comet/Asteroid Encounters: The STARDUST Example", AAS Paper 97-628, AAS/AIAA Astrodynamics Conference, Sun Valley, ID, August 1997.